Noname manuscript No.
(will be inserted by the editor)

Pipeline Implementations of Neumann–Neumann and Dirichlet–Neumann Waveform Relaxation Methods

Benjamin W. Ong · Bankim C. Mandal

arXiv:1605.08503v2 [math.NA] 5 Jan 2017

Received: January 6, 2017

Abstract This paper is concerned with the reformulation of Neumann-Neumann Waveform Relaxation (NNWR) methods and Dirichlet-Neumann Waveform Relaxation (DNWR) methods, a family of parallel space-time approaches to solving time-dependent PDEs. By changing the order of the operations, pipeline-parallel computation of the waveform iterates are possible without changing the final solution. The parallel efficiency and the increased communication cost of the pipeline implementation is presented, along with weak scaling studies to show the effectiveness of the pipeline NNWR and DNWR algorithms.

Keywords Dirichlet–Neumann; Neumann–Neumann; Waveform Relaxation; Domain Decomposition

Mathematics Subject Classification (2000) 65M55, 65Y05, 65M20

1 Introduction

Dirichlet–Neumann waveform relaxation (DNWR) and Neumann–Neumann waveform relaxation (NNWR) methods [8,10,19,14,9] have been formulated and analyzed recently for parabolic and hyperbolic partial differential equations (PDEs). These iterative methods are based on non-overlapping domain decomposition in space, where the iterations require subdomain solves with Dirichlet or Neumann boundary conditions. The superlinear convergence behavior of both the DNWR and NNWR methods have previously been shown for the heat equation [8,10]; finite step convergence to the exact solution for the wave equation has also been shown [9,20].

Both the DNWR and NNWR methods belong to the family of Waveform Relaxation (WR) methods. In the classical WR formulation, one solves the space-time

Benjamin W. Ong
Mathematical Sciences, Michigan Technological University, Houghton, MI, 49931
ongbw@mtu.edu

Bankim C. Mandal
Dept. of Mathematics, Michigan State University, East Lansing, MI, 48824
bmandal@math.msu.edu

subproblem over the entire time horizon at each iteration, before communicating interface data across subdomains. These WR methods have their origin in the work of Picard–Lindelöf [24,17] in the 19th century; WR methods were introduced as a parallel approach for solving systems of ODEs [16]. Since then, the WR framework has been combined with numerical approaches for solving elliptic problems for tackling time-dependent parabolic PDEs [11,13]. New transmission conditions have also been developed, known as optimized Schwarz WR (OSWR) methods, in order to achieve convergence with non-overlapping subdomains, or in general, faster convergence for overlapping subdomains; OSWR approaches for parabolic problems [6] and OSWR approaches for hyperbolic problems [7] have both been explored. This paper discusses a pipeline implementation of the DNWR and NNWR algorithms [18], which are extensions of Dirichlet-Neumann [1,3,21,22] and Neumann-Neumann [2,4,15,25] algorithms for solving space–time PDEs. By carefully re-arranging the order of operations, different waveform iterates can be simultaneously computed in a pipeline-parallel fashion, *without* changing the final solution.

The pipeline implementation of the DNWR and NNWR algorithm subdivides the entire time domain into smaller time blocks. Each space-time subproblem is solved on this smaller time block, and updated transmission conditions are transmitted before advancing to the next time block. This enables many concurrent subdomain space-time blocks, resulting in efficient parallelization on modern supercomputers. Specifically, given an appropriate number of processors, many waveform iterates can be computed in the same wall-clock time as a single processor computing one waveform iterate. The pipeline implementation for WR relaxation was mentioned in [26,12], and recently reintroduced and benchmarked for the classical Schwarz WR method [23].

The theoretical presentation of the algorithms described in this paper are for one-dimensional time-dependent PDEs of the form

$$\partial_t u - \mathcal{L}u = f(x,t), \quad (x,t) \in \Omega, \tag{1a}$$

$$u(x,0) = u_0(x), \quad x \in [0,L], \tag{1b}$$

$$u(0,t) = g_l(t), \quad u(L,t) = g_r(t), \quad t \in [0,T], \tag{1c}$$

where $\mathcal{L}$ is a spatial operator, the space-time domain, $\Omega : [0,L]\times[0,T]$, is a bounded domain, $\partial\Omega$ is a smooth boundary, $u_0(x)$ is the initial condition, and $g_l(t)$ and $g_r(t)$ are Dirichlet boundary conditions. The pipeline implementations can also be applied naturally to time-dependent PDEs of the form $\partial_{tt} - \mathcal{L}(u) = f(x,t)$, as well as WR methods formulated for solving PDEs in higher spatial dimensions.

This paper is broken into two main sections. In Section 2, we review the NNWR method for solving equation (1) before presenting the pipeline approach and numerical studies. The DNWR algorithm is outlined in Section 3, along with various pipeline approaches and numerical results.

2 Neumann–Neumann Waveform Relaxation (NNWR)

The NNWR algorithm for the model problem (1) with multi–subdomains setting was previously proposed and analyzed in [8]. In this method, the space-time domain $\Omega : [0,L] \times [0,T]$ is first partitioned into non-overlapping space-time subdomains, $\{\Omega_i,\ 1 \le i \le N\}$. Figure 1 illustrates a simple 1D spatial

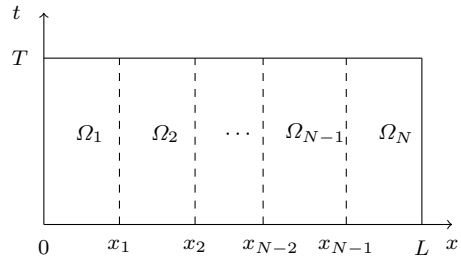

Fig. 1: Decomposition of a space-time domain, $[0, L] \times [0, T]$ domain

decomposition, although more complicated 2D or 3D decompositions (with cross points) are also possible. Let $\partial\Omega_i$ denote the domain boundary of Ω_i, and let $w_i^{[0]} = w(x_i, t), i = 1, \ldots, N-1$, be an initial time-dependent guess on the subdomain boundaries. The NNWR method performs a two-step iteration for each waveform iterate, $k = 1, 2, \ldots$, until convergence is reached. This two-step iteration consists of first solving a "Dirichlet subproblem" on each space-time subdomain,

$$\partial_t u_i^{[k]} - \mathcal{L} u_i^{[k]} = f, \quad (x,t) \in \Omega_i, \tag{2a}$$

$$u_i^{[k]}(x,0) = u_0(x), \quad x \in (x_{i-1}, x_i), \tag{2b}$$

$$u_i^{[k]}(x_{i-1},t) = \begin{cases} g_l(t) & \text{if } i = 1 \\ w_{i-1}^{[k-1]}(t) & \text{otherwise} \end{cases}, \tag{2c}$$

$$u_i^{[k]}(x_i,t) = \begin{cases} g_r(t) & \text{if } i = N \\ w_i^{[k-1]}(t) & \text{otherwise} \end{cases}, \tag{2d}$$

followed by solving an auxiliary "Neumann subproblem"

$$\partial_t \psi_i^{[k]} - \mathcal{L} \psi_i^{[k]} = 0, \quad (x,t) \in \Omega_i \tag{3a}$$

$$\psi_i^{[k]}(x,0) = 0, \quad x \in (x_{i-1}, x_i), \tag{3b}$$

$$\begin{cases} \psi_1^{[k]}(0) = 0, & \text{if } i = 1, \\ -\partial_x \psi_i^{[k]}(x_{i-1},t) = (\partial_x u_{i-1}^{[k]} - \partial_x u_i^{[k]})(x_{i-1},t), & \text{if } i > 1, \end{cases} \tag{3c}$$

$$\begin{cases} \partial_x \psi_i^{[k]}(x_i,t) = (\partial_x u_i^{[k]} - \partial_x u_{i+1}^{[k]})(x_i,t), & \text{for } i < N, \\ \psi_N^{[k]}(L) = 0, & \text{if } i = N. \end{cases} \tag{3d}$$

Then, the Dirichlet traces at the subdomain interfaces are updated,

$$w_i^{[k]}(t) = w_i^{[k-1]}(t) - \theta \left(\psi_i^{[k]}(x_i,t) + \psi_{i+1}^{[k]}(x_i,t) \right). \tag{4}$$

Remark 1 The auxiliary equations are only solved if the waveform iterates have not converged. The convergence of the NNWR algorithm for the heat equation was

previously analyzed [8]. For $\theta = 1/4$, the NNWR algorithm converges superlinearly with the estimate

$$\max_{1\le i\le N-1} \|w_i^{[k]}\|_{L^\infty(0,T)} \le \left(\frac{\sqrt{6}}{1-e^{-\frac{(2k+1)\tilde{h}^2}{T}}}\right)^{2k} e^{-k^2\tilde{h}^2/T} \max_{1\le i\le N-1} \|w_i^{[0]}\|_{L^\infty(0,T)},$$

where $\tilde{h} = \min\{x_i - x_{i-1} : i = 1, \ldots, N\}$ is the minimum subdomain width.

For the wave equation, the NNWR algorithm converges after the second waveform iterate, provided the window of integration is sufficiently small.

2.1 Classical NNWR implementation

A pseudo-code for a classical implementation of the NNWR in $\mathbb{R}^1$ is given in Algorithm 1. If the domain is broken into N non-overlapping subdomains, the classical implementation uses a total of N processing cores to approximate the solution. $2(N-1)(2K-1)$ or $4(N-1)K$ messages are needed (depending on when the stopping criterion is satisfied), with each message containing N_t words, where K is the number of waveform iterates computed and N_t is the number of time steps. In this simplified pseudo-code, it is assumed that identical time discretizations are taken in each subdomain. A generalization where each subdomain might take different discretizations is possible, provided an interpolation of the Neumann or Dirichlet traces are computed before lines 10 and 16 of Algorithm 1.

```
1  for i = 1 to N do in parallel
2      if i < N then
3          Guess w_i^[0](t_ℓ),   ℓ = 1, ..., N_t;
4      if i > 1 then
5          Guess w_{i-1}^[0](t_ℓ),   ℓ = 1, ..., N_t;
6      Set k = 1;
7      while not converged do
8          for ℓ = 1 to N_t do
9              Solve equation (2) for u_i^[k](x, t_ℓ);
10             Compute jump in Neumann data along interfaces;
11         Transmit/Receive Neumann data;
12         Check for convergence. If converged, break;
13         if not converged then
14             for ℓ = 1 to N_t do
15                 Solve equation (3) for ψ_i^[k](x, t_ℓ);
16                 Solve equation (4) for w_i^[k](t_ℓ);
17             Transmit/Receive Dirichlet data;
18             Check for convergence;
19         k ← k + 1;
```

Algorithm 1: The classical implementation of the NNWR algorithm is able to utilize N computing cores if the domain is broken into N non-overlapping subdomains.

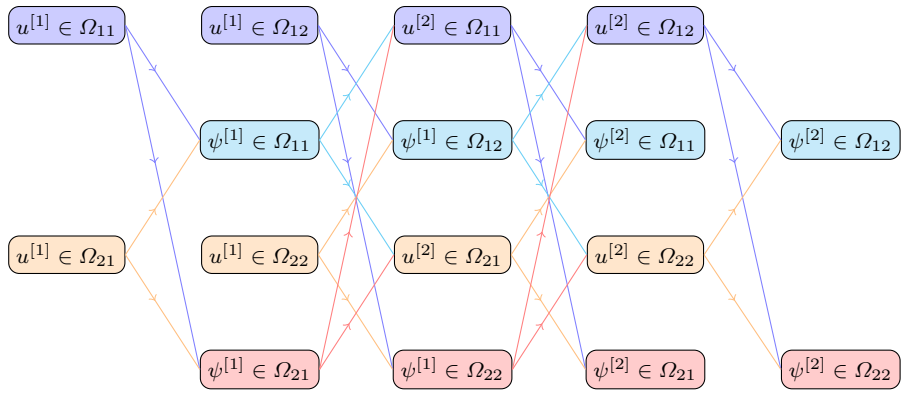

Fig. 2: A graphical illustration of a pipeline NNWR implementation for a two-domain, two time block example. The different color blocks represent different processing cores – in this case, four processors can be used to compute this pipeline NNWR implementation.

2.2 Pipeline NNWR implementation

A pipeline implementation of the NNWR algorithm allows for higher concurrency at the expense of increased communication: multiple waveform iterates are simultaneously evaluated after an initial start-up cost. Both implementations will return the same numerical solution. The main idea is to decompose the time window into non-overlapping blocks, and transmit messages to available processors after the evaluation of each time block, so that other processors can simultaneously compute solutions to the auxiliary equations or the next waveform iterate. We first illustrate this for a simple two subdomains, two time blocks example. Let $0 = T_0 < T_1 < T_2 = T$ and let Ω_{ij} denotes the space-time domain, $[x_{i-1}, x_i] \times [T_{j-1}, T_j]$. The pipeline implementation starts with solving the Dirichlet subproblem, equations (2), for $u^{[1]}$ in Ω_{11} and Ω_{21} in parallel with two processors. As soon as this computation is completed, messages are sent to available processors. The algorithm then proceeds with solving the Dirichlet subproblems for $u^{[1]}$ in Ω_{12} and Ω_{22}, as well as solving auxiliary equation (3) for $\psi^{[1]}$ in Ω_{11} and Ω_{21} simultaneously, using a total of four processors. If the waveform iterates have not converged, messages are sent to the appropriate processors, and the algorithm proceeds with solving the Dirichlet subproblems for $u^{[2]}$ in Ω_{11} and Ω_{21}, as well as solving the auxiliary subproblem for $\psi^{[1]}$ in Ω_{12} and Ω_{22}, again with four processors. A graphical illustration of the pipeline framework for this example, assuming two full waveform iterates are desired, is shown in Figure 2.

A pseudo-code for a pipeline implementation of the NNWR in $\mathbb{R}^1$ for the case $J \geq 2K$ is given in Algorithm 2. If the domain is broken into $N \times J$ non-overlapping subdomains as shown in Figure 3, this implementation is able to utilize up to $2NK$ processing cores, where K is the number of waveform iterates computed and J is the number of time blocks. The number of messages increases, with up to $4J(N-1)K$ messages needed in the algorithm. The size of each message decreases to N_t/J words. If $J < 2K$, the pipeline implementation is only able to utilize

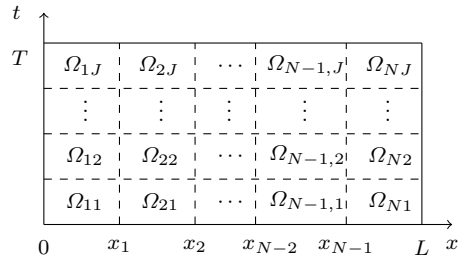

Fig. 3: Decomposition of a space-time domain, $[0, L]\times[0, T]$, into $N\times J$ subdomains for pipeline parallelism.

NJ processors. A pseudo-code for a pipeline NNWR implementation for the case $J < 2K$ is given in Algorithm 6 in the Appendix.

Remark 2 Unlike the classical NNWR implementation where one iterates until convergence, Algorithm 2 requires a specification of K, the number of waveform iterates to be computed. One can use a priori error estimates to pick K intelligently, but it is likely that there will either be wasted work (if convergence is obtained for $k < K$) or an unconverged solution, if K is not large enough. All is not lost however if K is not large enough, since the computation can be restarted using the most accurate Dirichlet traces. Also, in support of dynamic resource allocation and fault resiliency, there is active development to allow MPI communicators to expand or shrink. However, there are no plans for adopting dynamic MPI communicators into MPI standards in the near future.

Remark 3 The pipeline NNWR implementation has a start-up and shut-down phase before multiple waveform iterates can be computed in parallel, i.e., at the start and end of the computation, some processing cores sit idle. For K full waveform iterates and J time blocks, the peak theoretical parallel efficiency is

$$\begin{cases} \frac{2K}{2K+J-1}, & \text{if} \quad 2K \geq J, \\ \frac{J}{2K+J-1}, & \text{if} \quad 2K < J. \end{cases} \tag{5}$$

2.3 Numerical Experiments

As mentioned in the introduction, the discretization and convergence of NNWR algorithms have already been discussed and established for parabolic and hyperbolic PDEs [9,8,18]. This section is concerned with the efficacy of the pipeline implementation, where many waveform iterates can be concurrently computed. The heat equation is solved, $u_t = u_{xx}$, subject to the initial conditions $u(x, 0) = (x - 0.5)^2 - 0.25$ and homogeneous Dirichlet boundary data. The computational domain, $\Omega \times [0, T]$, is $[0, 1] \times [0, 0.1]$. A centered finite difference is used to approximate the spatial operator, and a backward Euler integration is used. Each subdomain solves the linear system for each time advance by performing a backwards and forwards substitution using a pre-factored LU decomposition of the corresponding matrix. The `C` code, available on the author's website, uses MPI

```
   /* Assumes J ≥ 2K                                                    */
 1 for i = 1 to N do in parallel
 2     for k = 1 to K do in parallel
 3         for q = 0 to 1 do in parallel
 4             if q == 1 then
                   /* Dirichlet Update                                  */
 5                 for j = 1 to J do
 6                     if k > 1 then
 7                         receive updated Dirichlet data;
 8                     else
 9                         for ℓ = 1 to N_t/J do
10                             t_ℓ = (ℓ + N_t(j−1)/J) Δt;
11                             if i < N then
12                                 Guess w_i^[0](t_ℓ);
13                             if i > 1 then
14                                 Guess w_{i−1}^[0](t_ℓ);
15                         for ℓ = 1 to N_t/J do
16                             Solve equation (2) for u_i^[k](x, t_ℓ);
17                             Compute jump in Neumann data along interfaces;
18                         Send Neumann data;
19                     Check for convergence. If converged, break;
20             else
                   /* Auxiliary update                                  */
21                 for j = 1 to J do
22                     Receive Neumann data;
23                     for ℓ = 1 to N_t/J do
24                         t_ℓ = (ℓ + N_t(j−1)/J) Δt;
25                         Solve equation (3) for ψ_i^[k](x, t_ℓ);
26                         Solve for w_i^[k](t_ℓ) using equation (4);
27                     if k < K then
28                         Send Dirichlet data;
29                     Check for convergence. If converged, break;
```

Algorithm 2: This pipeline implementation of the NNWR algorithm is able to utilize $2NK$ computing cores if the domain is broken into N non-overlapping subdomains and K full iterates are used, provided $J \geq 2K$.

parallelism. The reported numerical experiments were performed using the Stampede supercomputer at the Texas Advanced Computing Center. The resources were provided through the NSF-supported Extreme Science and Engineering Discovery Environment (XSEDE) program.

The first numerical experiment validates the qualitative behavior of the theoretical peak efficiency, equation (5), as the number of time blocks, J, is varied. The number of full waveform iterates is fixed at $K = 4$, the number of subdomains is fixed at $N = 8$, the spatial and temporal discretizations are fixed at $N_x = 32000$, $N_t = 8192$. A total of 64 processing cores are used in this first experiment. Figure 4 displays the expected behavior – for a small number of time blocks, processors sit idle for a larger percentage of time, leading to poor efficiency.

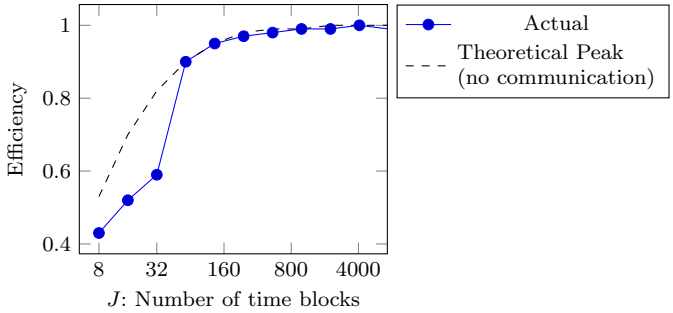

Fig. 4: Efficiency of the pipeline NNWR implementation as a function of the number of time blocks, J. For a large number of time blocks, J, the algorithm is able to utilize all processors in a pipeline fashion for a larger percentage of the computation, leading to higher efficiency.

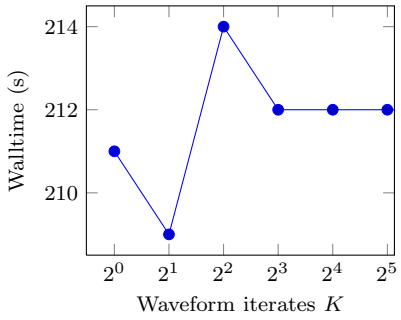

Fig. 5: Weak Scaling for Pipeline NNWR: Wall time vs Waveform iterations. The pipeline implementation scales weakly with almost no overhead, i.e., with $2NK$ processing cores, we can compute K iterations in almost the same walltime as computing one iteration with $2N$ processing cores.

For large number of time blocks, the pipeline implementation attains close to theoretical peak efficiency, indicating that the communication overhead is negligible. Here, an efficiency close to 1 means that the pipeline NNWR implementation with $2KN$ processing cores is able to compute K full waveform iterates $2K$ times faster than the the classical NNWR implementation using N processors. The data used to generate Figure 4 is summarized in Appendix, Table A1.

A weak scaling study is performed for the second numerical experiment. For a fixed spatial, temporal and time-block discretization, the number of processors is varied to compute a differing number of waveform iterates. In the experiment, we pick $N = 8$ subdomains with $N_x = 32000$ and $N_t = 8192$, $J = 1024$ time blocks and $2KN$ processor cores. Figure 5 shows that the pipeline NNWR implementation is able to scale weakly with very good efficiency. The data used to generate Figure 5 is summarized in the Appendix, Table A2.

3 Dirichlet–Neumann Waveform Relaxation (DNWR)

The DNWR method [8,10] can be implemented in various arrangements, in terms of how the Dirichlet and Neumann transmission conditions are enforced along artificial boundaries. Pipeline parallelism can be applied to all the variants, although the efficacy will vary with each variant. This section discusses the pipeline implementation applied to the arrangement proposed by Funaro, Quateroni and Zanolli [5]. Using the notation for the subdomain discretization in Section 2, the DNWR algorithm also requires an initial guess for the solution on the subdomain boundaries, $w_i^{[0]}(t), i = 1, \dots N$, as well as a selection of the initial subdomain, Ω_m, with $1 \leq m \leq N$; picking m to be the middle subdomain, e.g., $m = \lceil N/2 \rceil$ will reduce the start-up overhead for the DNWR algorithm. Each space-time subdomain solves one of the following subproblems. In subdomain m, a Dirichlet subproblem is solved,

$$\partial_t u_m^{[k]} - \mathcal{L} u_m^{[k]} = f, \quad (x,t) \in \Omega_m, \tag{6a}$$

$$u_m^{[k]}(x,0) = u_0(x), \quad x \in (x_{m-1}, x_m), \tag{6b}$$

$$u_m^{[k]}(x_{m-1},t) = \begin{cases} g_l(t) & \text{if } m = 1 \\ w_{m-1}^{[k-1]}(t) & \text{otherwise} \end{cases}, \tag{6c}$$

$$u_m^{[k]}(x_m,t) = \begin{cases} g_r(t) & \text{if } m = N \\ w_m^{[k-1]}(t) & \text{otherwise} \end{cases}, \tag{6d}$$

For subdomains to the left of subdomain m, i.e., $i < m$, a Dirichlet-Neumann subproblem is solved,

$$\partial_t u_i^{[k]} - \mathcal{L} u_i^{[k]} = f, \quad (x,t) \in \Omega_i, \tag{7a}$$

$$u_i^{[k]}(x,0) = u_0(x), \quad x \in (x_{i-1}, x_i), \tag{7b}$$

$$u_i^{[k]}(x_{i-1},t) = \begin{cases} g_l(t) & \text{if } i = 1 \\ w_{i-1}^{[k-1]}(t) & \text{otherwise} \end{cases}, \tag{7c}$$

$$\partial_x u_i^{[k]}(x_i,t) = \partial_x u_{i+1}^{[k]}(x_i,t) \tag{7d}$$

For subdomains to the right of subdomain m, i.e., $i > m$, a Neumann–Dirichlet subproblem is solved,

$$\partial_t u_i^{[k]} - \mathcal{L} u_i^{[k]} = f, \quad (x,t) \in \Omega_i, \tag{8a}$$

$$u_i^{[k]}(x,0) = u_0(x), \quad x \in (x_{i-1}, x_i), \tag{8b}$$

$$\partial_x u_i^{[k]}(x_{i-1},t) = \partial_x u_{i-1}^{[k]}(x_{i-1},t) \tag{8c}$$

$$u_i^{[k]}(x_i,t) = \begin{cases} g_r(t) & \text{if } i = N \\ w_i^{[k-1]}(t) & \text{otherwise} \end{cases}, \tag{8d}$$

After the subdomain problems are solved, the Dirichlet traces are updated,

$$w_i^{[k]}(t) = \theta u_i^{[k]}(x_i,t) + (1-\theta) w_i^{[k-1]}(t), \quad i < m, \tag{9a}$$

$$w_i^{[k]}(t) = \theta u_{i+1}^{[k]}(x_i,t) + (1-\theta) w_i^{[k-1]}(t), \quad i \geq m. \tag{9b}$$

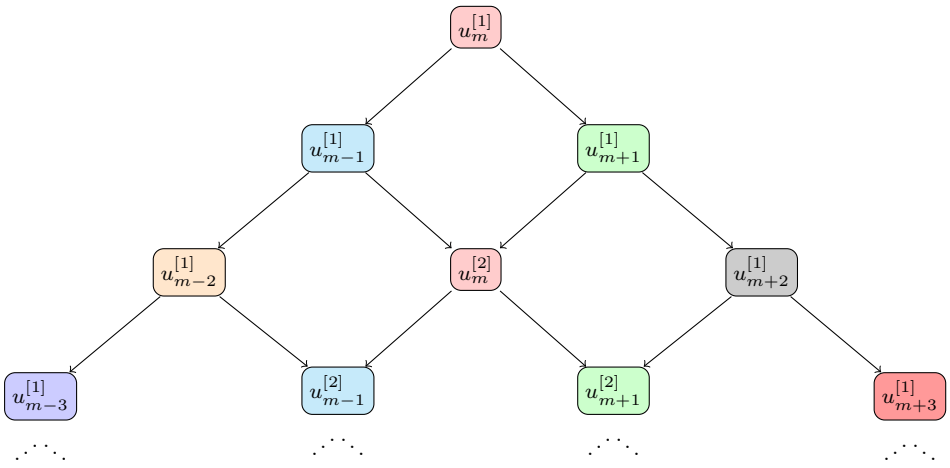

Fig. 6: DNWR implemented classically. The different colors represent N different processing cores (inefficiently) used to compute the waveform iterates in different space–time subdomains.

Remark 4 The DNWR algorithm converges and the rates of convergence have been analyzed for $N = 2$ in [8] and for $N > 2$ in [10]. For the case $N > 2$, $\theta = 1/2$ and $m = \lceil N/2 \rceil$, the DNWR decomposition of the heat equation in $\mathbb{R}^1$ results in the following superlinear convergence estimate:

$$\max_{1\le i\le N-1} \|w_i^{[k]}\|_{L^\infty(0,T)} \le \left(N - 4 + \frac{2h_{\max}}{h_m}\right)^k \operatorname{erfc}\left(\frac{kh_{\min}}{2\sqrt{T}}\right) \max_{1\le i\le N-1} \|w_i^{[0]}\|_{L^\infty(0,T)},$$

where $h_i = x_i - x_{i-1}$, $h_{\max} := \max\limits_{1\le i\le N} h_i$ and $h_{\min} := \min\limits_{1\le i\le N} h_i$.

3.1 Classical DNWR implementation

Equations (6) – (7) – (8) – (9) can be "decoupled" by solving the space–time subproblems in a specific order. Recall the decomposition of the space-time domain as described in Figure 1. One starts by computing $u_m^{[1]}(x,t)$ in Ω_m and transmitting the computed boundary conditions before computing $u_{m-1}^{[1]}(x,t)$ and $u_{m+1}^{[1]}(x,t)$ in the neighboring subdomains. The updated boundary conditions are transmitted, and then $u_{m-2}^{[1]}(x,t)$, $u_m^{[2]}(x,t)$ and $u_{m+1}^{[1]}(x,t)$ are simultaneously computed. If optimal parallel efficiency is not a concern, N processes can be spawned as shown in Algorithm 3. Not all processes can be utilized simultaneously however, because transmission conditions needed in lines 6 through 19 may not be available until other processes have computed and transmitted the required information. Although not a practical algorithm, this pseudo code is instructive because it outlines the flow of information in an easy to read fashion. Figure 6 shows the flow of information for the first few steps of the DNWR algorithm, implemented classically using N processors and N subdomains.

```
1  for i = 1 to N do in parallel
2      m = nprocs = ⌈N/2⌉;
3      if m ≤ i < N then
4          Guess w_i^[0](t_ℓ),   ℓ = 1, ... N_t ;
5      if 1 < i ≤ m then
6          Guess w_{i-1}^[0](t_ℓ),   ℓ = 1, ... N_t ;
7      for k = 1 to K do
           /* receive boundary information                    */
8          switch i do
9              case i = m
10                 if k > 1 then
11                     Receive Dirichlet data from neighbors (if they exist);
12                     Update Dirichlet traces using equation (9);
13             case i < m
14                 Receive Neumann data from right neighbor;
15                 if k > 1 and i > 1 then
16                     Receive Dirichlet data from left neighbors;
17                     Update Dirichlet traces using equation (9);
18             case i > m
19                 Receive Neumann data from left neighbor;
20                 if k > 1 and i < N then
21                     Receive Dirichlet data from right neighbors;
22                     Update Dirichlet traces using equation (9);
           /* Solve space time problem                        */
23         for ℓ = 1 to N_t do
24             switch i do
25                 case i = m
26                     Solve equation (6) for u_i^[k](x, t_ℓ);
27                 case i < m
28                     Solve equation (8) for u_i^[k](x, t_ℓ);
29                 case i > m
30                     Solve equation (7) for u_i^[k](x, t_ℓ);
           /* Send boundary conditions                         */
31         switch i do
32             case i = m
33                 Send Neumann data to neighbors (if they exist);
34             case i < m
35                 if i > 1 then
36                     Send Neumann/convergence data to left neighbor;
37                 if k < K then
38                     Send Dirichlet data to right neighbor;
39             case i > m
40                 if i < N then
41                     Send Neumann/convergence data to right neighbor;
42                 if k < K then
43                     Send Dirichlet data to left neighbor;
44         Check for convergence. If converged, break;
```

Algorithm 3: A classical (naive) implementation of the DNWR algorithm using N processes, where K is the number of waveform iterates and N is the number of non-overlapping spatial subdomains.

Remark 5 By sending appropriate convergence flags in lines 37 and 42 of Algorithm 3, the processors computing the solution to the space time block in Ω_1 and Ω_N are able to determine if the algorithm has converged.

To construct practical DNWR algorithms using a classical implementation, recall from the previous section that selecting $m = \lceil N/2 \rceil$ reduces the startup overhead for the DNWR algorithm, thereby improving parallel efficacy. For the remainder of this paper, assume that $m = \lceil N/2 \rceil$. The classical DNWR algorithm depends on the ratio between the number of subdomains and the number of waveform iterates being computed. If $2K \geq \lceil N/2 \rceil$, the DNWR algorithm is able to utilize $\lceil N/2 \rceil$ processing cores – each processor computes the solution for two spatial subdomains. Algorithm 4 outlines a pseudo code of the DNWR implemented classically. Algorithm 7 in the Appendix discusses the implementation for the case $2K < \lceil N/2 \rceil$. For both cases, a total of $(N-1)(2K-1)$ messages are needed, with each message containing N_t words, where N_t is the number of time steps.

Remark 6 Unlike the classical implementation of the NNWR algorithm, the classical implementation of the DNWR algorithm requires specification of K, the number of waveform iterates desired. This is because each spatial subdomain might compute different waveform iterates simultaneously. Similar to Remark 2, a priori specification of K might result in an unconverged solution if the prescribed K is not large enough, or wasted work if convergence is obtained for $k < K$ iterations.

3.2 Pipeline DNWR implementation

Similar to the pipeline NNWR implementation in Section 2.2, a pipeline implementation of the DNWR algorithm allows for higher concurrency at the expense of increased communication: multiple waveform iterates are simultaneously evaluated (in $\cup_j \Omega_{ij}$) after an initial startup cost. As before, both the classical implementation and the pipeline implementation of the DNWR algorithm will result in the same numerical solution. Figure 7 shows a graphical representation of a five-domain, two waveform iterate example. For a general decomposition as shown in Figure 3, the pipeline DNWR implementation is outlined in Algorithm 5. If the number of time blocks, J, are sufficiently large, the pipeline DNWR algorithm is able to utilize NK processing cores. The pipeline DNWR implementation requires a total of $J(N-1)(2K-1)$ messages, with each message containing N_t/J words. Since the numerical results in Section 2.3 indicate that the communication overhead of utilizing many time blocks J is negligible, we restrict our discussion to the case where J is sufficiently large, $J > \lceil N/2 \rceil + (2K-1)$.

```
   Input: N: # subdomains; K: # waveform iterates
 1 m = ⌈N/2⌉.;
 2 for p = 1 to nprocs do in parallel
 3     Set i = 2(p − 1);
 4     if i > 0 && i ≤ m then
 5         Guess w_i^[0](t_ℓ),   ℓ = 1 . . . , N_t ;
 6     Set i = 2p;
 7     if i < N && i ≥ m then
 8         Guess w_i^[0](t_ℓ),   ℓ = 1 . . . , N_t ;
 9     Set i = 2p − 1;
10     if i < N then
11         Guess w_i^[0](t_ℓ),   ℓ = 1 . . . , N_t ;
12     for k = 1 to K do
13         switch p do
14             case 2p < m
15                 Set i = 2p;
16             case (2p − 1) > m
17                 Set i = 2p − 1;
18             case 2p = m
19                 Set i = 2p;
20             otherwise
21                 Set i = 2p − 1;
22         if i ∈ {1, . . . , N} then
               /* receive boundary info: lines 8--22 in Alg. 3        */
               /* Solve space time pde: lines 23--30 in Alg. 3        */
               /* Send boundary info: lines 31--43 in Alg. 3          */
23         switch p do
24             case 2p < m
25                 Set i = 2p − 1;
26             case (2p − 1) > m
27                 Set i = 2p;
28             case 2p = m
29                 Set i = 2p − 1;
30             otherwise
31                 Set i = 2p;
32         if i ∈ {1, . . . , N} then
               /* receive boundary info: lines 8--22 in Alg. 3        */
               /* Solve space time pde: lines 23--30 in Alg. 3        */
               /* Send boundary info: lines 31--43 in Alg. 3          */
33         Check for convergence. If converged, break;
```

Algorithm 4: A classical implementation of the DNWR algorithm using $\lceil N/2 \rceil$ processes, where N is the number of non-overlapping spatial subdomains.

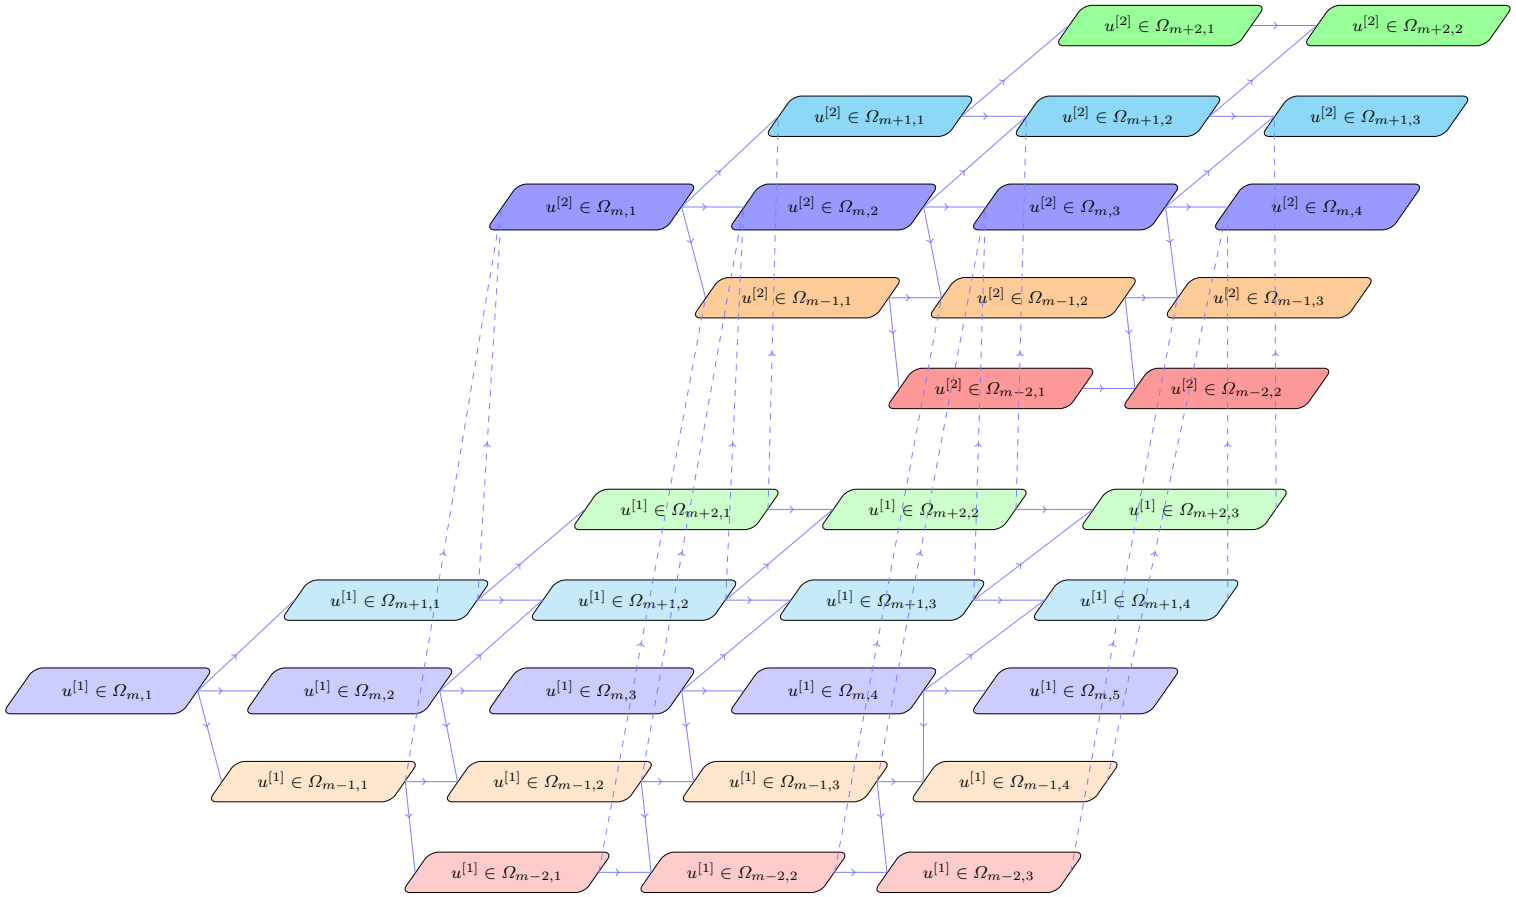

Fig. 7: A pipeline DNWR implementation for a five domain, two waveform iterate example. Each color/shade represents a different processor.

```
   /* Assumes J > ⌈N/2⌉ + (2K − 1)                                          */
   Input: N: # subdomains; K: # waveform iterates
 1 m = ⌈N/2⌉;
 2 for i = 1 to N do in parallel
 3     for k = 1 to K do in parallel
 4         for j = 1 to J do
               /* Receive boundary data, lines 8--22 in Alg. 3             */
 5             for ℓ = 1 to N_t/J do
 6                 t_ℓ = (ℓ + N_t(j−1)/J) Δt;
 7                 if k = 1 then
 8                     if i < N then
 9                         Guess w_i^[0](t_n);
10                     if i > 0 then
11                         Guess w_{i−1}^[0](t_n);
                   /* Solve space-time PDE, lines 23--30, Alg. 3           */
               /* Send boundary conditions, lines 31--43 in Alg. 3         */
12         Check for convergence. If converged, break;
```

Algorithm 5: The pipeline implementation of the DNWR algorithm is able to utilize NK computing cores if the domain is broken into N non-overlapping subdomains.

Remark 7 The pipeline DNWR implementation has a start-up and shut-down phase before multiple waveform iterates can be computed in parallel. For K waveform iterates and J time blocks, the parallel efficiency is

$$\frac{J}{J + \lfloor (N/2) \rfloor + 2(K-1)}, \tag{10}$$

provided $J > \lceil N/2 \rceil + (2K-1)$.

3.3 Numerical Experiment

A weak scaling study of the pipeline DNWR algorithm concludes the discussion of implementation details related to waveform relaxation methods. The heat equation described in Section 2.3 is solved with a fixed spatial, temporal and time-block discretization. The number of processors is varied to compute a differing number of waveform iterates. In the experiment, we pick $N = 8$ subdomains with $N_x = 32,000$, $N_t = 8192$ and $J = 1024$ time blocks. Figure 8 shows that the pipeline DNWR implementation is able to scale weakly with very good efficiency. The data used to generate Figure 8 is summarized in the Appendix, Table A3.

4 Conclusions

In this paper, we have reformulated the NNWR and DNWR methods to allow for pipeline-parallel computation of the waveform iterates, after an initial start-up cost. The key observation is that the order of the computations (loops) can be

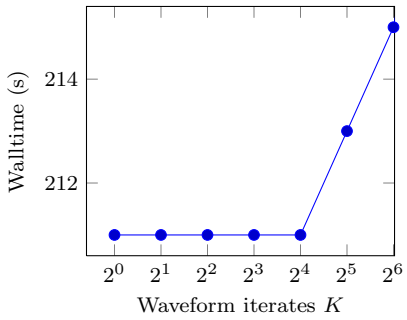

Fig. 8: Weak Scaling for Pipeline DNWR: Wall time vs Waveform iterations. The pipeline implementation scales weakly with almost no overhead, i.e., with NK processing cores, we can compute K iterations in almost the same walltime as computing one iteration with N processing cores.

flipped without affecting the final solution, at the expense of increased communication. Theoretical estimates for the parallel speedup and communication overhead are presented, along with weak scaling studies to show the effectiveness of the pipeline DNWR and NNWR algorithms.

A significant drawback, is that the pipeline implementations (and the DNWR algorithm) require an apriori specification of the number of waveform iterates to be computed. Hence, it is likely that there will either be wasted work (if convergence is obtained for fewer waveform iterates) or an unconverged solution, if insufficient waveform iterates were specified. An adaptive approach is being designed and analyzed to improve these pipeline parallel computations.

Acknowledgments

This work utilized computational resources provided by superior, the high-performance computing cluster @ MTU, and by the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1053575.

References

1. Bjørstad, P.E., Widlund, O.B.: Iterative methods for the solution of elliptic problems on regions partitioned into substructures. SIAM J. Numer. Anal. **23**(6), 1097–1120 (1986). DOI 10.1137/0723075. URL http://dx.doi.org/10.1137/0723075
2. Bourgat, J.F., Glowinski, R., Le Tallec, P., Vidrascu, M.: Variational formulation and algorithm for trace operator in domain decomposition calculations. In: Domain decomposition methods (Los Angeles, CA, 1988), pp. 3–16. SIAM, Philadelphia, PA (1989)
3. Bramble, J.H., Pasciak, J.E., Schatz, A.H.: An iterative method for elliptic problems on regions partitioned into substructures. Math. Comp. **46**(174), 361–369 (1986). DOI 10.2307/2007981. URL http://dx.doi.org/10.2307/2007981
4. De Roeck, Y.H., Le Tallec, P.: Analysis and test of a local domain-decomposition preconditioner. In: Fourth International Symposium on Domain Decomposition Methods for Partial Differential Equations (Moscow, 1990), pp. 112–128. SIAM, Philadelphia, PA (1991)

5. Funaro, D., Quarteroni, A., Zanolli, P.: An iterative procedure with interface relaxation for domain decomposition methods. SIAM J. Numer. Anal. **25**(6), 1213–1236 (1988). DOI 10.1137/0725069. URL http://dx.doi.org/10.1137/0725069
6. Gander, M.J., Halpern, L.: Optimized Schwarz waveform relaxation methods for advection reaction diffusion problems. SIAM J. Numer. Anal. **45**(2), 666–697 (electronic) (2007). DOI 10.1137/050642137. URL http://dx.doi.org/10.1137/050642137
7. Gander, M.J., Halpern, L., Nataf, F.: Optimal Schwarz waveform relaxation for the one dimensional wave equation. SIAM J. Numer. Anal. **41**(5), 1643–1681 (2003). DOI 10.1137/S003614290139559X. URL http://dx.doi.org/10.1137/S003614290139559X
8. Gander, M.J., Kwok, F., Mandal, B.C.: Dirichlet-Neumann and Neumann-Neumann Waveform Relaxation Algorithms for Parabolic Problems. Electron. Trans. Numer. Anal. **45**, 424–456 (2016)
9. Gander, M.J., Kwok, F., Mandal, B.C.: Dirichlet-Neumann and Neumann-Neumann Waveform Relaxation for the Wave Equation. In: T. Dickopf, M.J. Gander, L. Halpern, R. Krause, L. Pavarino (eds.) Domain Decomposition Methods in Science and Engineering XXII, Lecture Notes in Computational Science and Engineering, pp. 501–509. Springer International Publishing (2016). DOI 10.1007/978-3-319-18827-0_51. URL http://dx.doi.org/10.1007/978-3-319-18827-0_51
10. Gander, M.J., Kwok, F., Mandal, B.C.: Dirichlet-Neumann Waveform Relaxation Method for the 1D and 2D Heat and Wave Equations in Multiple subdomains. to appear (arXiv:1507.04011)
11. Gander, M.J., Stuart, A.M.: Space-time continuous analysis of waveform relaxation for the heat equation. SIAM J. Sci. Comput. **19**(6), 2014–2031 (1998). DOI 10.1137/S1064827596305337. URL http://dx.doi.org/10.1137/S1064827596305337
12. Gear, C.W.: Waveform methods for space and time parallelism. In: Proceedings of the International Symposium on Computational Mathematics (Matsuyama, 1990), vol. 38, pp. 137–147 (1991). DOI 10.1016/0377-0427(91)90166-H. URL http://dx.doi.org/10.1016/0377-0427(91)90166-H
13. Giladi, E., Keller, H.: Space time domain decomposition for parabolic problems. Tech. Rep. 97-4, Center for research on parallel computation CRPC, Caltech (1997)
14. Kwok, F.: Neumann-Neumann Waveform Relaxation for the Time-Dependent Heat Equation. In: J. Erhel, M.J. Gander, L. Halpern, G. Pichot, T. Sassi, O.B. Widlund (eds.) Domain Decomposition in Science and Engineering XXI, vol. 98, pp. 189–198. Springer-Verlag (2014)
15. Le Tallec, P., De Roeck, Y.H., Vidrascu, M.: Domain decomposition methods for large linearly elliptic three-dimensional problems. J. Comput. Appl. Math. **34**(1), 93–117 (1991). DOI 10.1016/0377-0427(91)90150-I. URL http://dx.doi.org/10.1016/0377-0427(91)90150-I
16. Lelarasmee, E., Ruehli, A., Sangiovanni-Vincentelli, A.: The waveform relaxation method for time-domain analysis of large scale integrated circuits. IEEE Trans. Compt.-Aided Design Integr. Circuits Syst. **1**(3), 131–145 (1982)
17. Lindelöf, E.: Sur l'application des méthodes d'approximations successives à l'étude des intégrales réelles des équations différentielles ordinaires. Journal de Mathématiques Pures et Appliquées pp. 117–128 (1894)
18. Mandal, B.C.: Convergence Analysis of Substructuring Waveform Relaxation Methods for Space-time Problems and Their Application to Optimal Control Problems. Ph.D. thesis, University of Geneva (2014). URL http://archive-ouverte.unige.ch/unige:46146
19. Mandal, B.C.: A Time-Dependent Dirichlet-Neumann Method for the Heat Equation. In: J. Erhel, M.J. Gander, L. Halpern, G. Pichot, T. Sassi, O.B. Widlund (eds.) Domain Decomposition in Science and Engineering XXI, vol. 98, pp. 467–475. Springer-Verlag (2014)
20. Mandal, B.C.: Neumann-Neumann Waveform Relaxation Algorithm in Multiple Subdomains for Hyperbolic Problems in 1d and 2d. Numer. Methods Partial Differ. Equ. (2016). URL DOI10.1002/num.22112
21. Marini, L.D., Quarteroni, A.: A relaxation procedure for domain decomposition methods using finite elements. Numer. Math. **55**(5), 575–598 (1989). DOI 10.1007/BF01398917. URL http://dx.doi.org/10.1007/BF01398917
22. Martini, L., Quarteroni, A.: An Iterative Procedure for Domain Decomposition Methods: a Finite Element Approach. SIAM, in Domain Decomposition Methods for PDEs, I pp. 129–143 (1988)

23. Ong, B.W., High, S., Kwok, F.: Pipeline schwarz waveform relaxation. In: T. Dickopf, M.J. Gander, L. Halpern, R. Krause, L. Pavarino (eds.) Domain Decomposition Methods in Science and Engineering XXII, Lecture Notes in Computational Science and Engineering, pp. 179–187. Springer International Publishing (2016). DOI 10.1007/978-3-319-18827-0_36. URL http://dx.doi.org/10.1007/978-3-319-18827-0_36
24. Picard, E.: Sur l'application des méthodes d'approximations successives à l'étude de certaines équations différentielles ordinaires. Journal de Mathématiques Pures et Appliquées pp. 217–272 (1893)
25. Toselli, A., Widlund, O.: Domain decomposition methods—algorithms and theory, *Springer Series in Computational Mathematics*, vol. 34. Springer-Verlag, Berlin (2005). DOI 10.1007/b137868. URL http://dx.doi.org/10.1007/b137868
26. Vandewalle, S.G., Van de Velde, E.F.: Space-time concurrent multigrid waveform relaxation. Ann. Numer. Math. **1**(1-4), 347–360 (1994). Scientific computation and differential equations (Auckland, 1993)

Appendix

J	Walltime (s)	Speedup	Actual Parallel Efficiency	Theoretical Peak Efficiency	MFLOPS/μs
8	492	6.82	0.43	0.53	46
16	402	8.35	0.52	0.70	57
32	357	9.40	0.59	0.82	64
64	233	14.40	0.90	0.90	98
128	221	15.18	0.95	0.95	103
256	216	15.53	0.97	0.98	105
512	213	15.75	0.98	0.99	107
1024	212	15.83	0.99	0.99	107
2048	211	15.90	0.99	1.00	108
4096	210	15.98	1.00	1.00	108
8192	211	15.90	0.99	1.00	108

Table A1: Raw data used to compute the efficiency of the pipeline NNWR implementation as a function of the number of time blocks, J. To compute the speedup and efficiency, the classical NNWR implementation with 8 processing cores was used – this benchmark computation took 3355 seconds.

```
   /* Assumes J < 2K                                                    */
 1 for i = 1 to N do in parallel
 2     niter = ⌈2K/J⌉;
 3     for iter = 1 to niter do
 4         for l = 1 to J do in parallel
 5             a = (iter − 1)J + l;
 6             if a ≤ 2K then
 7                 k = ⌈a/2⌉;
 8                 q = mod(a, 2);
 9                 if q == 1 then
                       /* Dirichlet Step: lines 5--19 in Algorithm 2   */
10                 else
                       /* Auxiliary Step: lines 21--29 in Algorithm 2  */
```

Algorithm 6: This pipeline implementation of the NNWR algorithm is able to utilize NJ computing cores if the domain is broken into N non-overlapping subdomains and K full iterates are used, provided $J < 2K$.

```
   /* assumes 2K ≤ ⌈N/2⌉                                                */
   Input: N: # subdomains; K: # waveform iterates
 1 m = ⌈N/2⌉;
 2 for p = 1 to 2K do in parallel
 3     niter = ⌈N/(4K)⌉;
 4     for iter = 1 to niter do
 5         for k = 1 to K do
 6             if p > K then
 7                 i = m + (iter − 1) ∗ 2K + 2(p − K − 1)
 8             else
 9                 i = m + (iter − 1) ∗ 2K + 2(p − K − 1) + 1
10             if i ∈ [1, N] then
                   /* receive boundary info: lines 9--23 in Alg. 3     */
                   /* Solve space time pde: lines 24--31 in Alg. 3     */
                   /* Send boundary info: lines  32--44 in Alg. 3      */
11             if p > K then
12                 i = m + (iter − 1) ∗ 2K + 2(p − K − 1) + 1
13             else
14                 i = m + (iter − 1) ∗ 2K + 2(p − K − 1)
15             if i ∈ [1, N] then
                   /* receive boundary info: lines 9--23 in Alg. 3     */
                   /* Solve space time pde: lines 24--31 in Alg. 3     */
                   /* Send boundary info: lines  32--44 in Alg. 3      */
16             Check for convergence. If converged, break;
```

Algorithm 7: A classical implementation of the DNWR algorithm using $2K$ processes, where K is the number of waveform iterates being computed. This pseudo code is for the case $2K \leq \lceil N/2 \rceil$.

K	# procs	Walltime (s)	Parallel Efficiency	MFLOPS/μs
1	16	211	1.00	27
2	32	209	1.01	54
4	64	214	0.99	108
8	128	212	1.00	213
16	256	212	1.00	429
32	512	212	1.00	815
64	1024	212	1.00	1540

Table A2: Raw data used to compute the weak scaling capability of the pipeline NNWR implementation. To compute the efficiency, the classical NNWR implementation with 8 processing cores compute one full iteration using 423 seconds. Here, an efficiency of 1 means that the pipeline NNWR implementation is able to compute K full waveform iterates using $2NK$ processing cores in half the time it takes N processing cores to compute one full waveform iterate.

K	# procs	Walltime (s)	MFLOPS/μs
1	8	211	3.4
2	16	211	5.0
4	32	211	8.4
8	64	211	15.1
16	128	211	28.5
32	256	212	54.8
64	512	215	106.5
128	1024	220	206.8

Table A3: Raw data used to compute the weak scaling capability of the pipeline DNWR implementation. The pipeline DNWR allows for multiple waveform iterates to be computed in approximately the same wall-clock time as eight processors computing a single waveform iterate.